

06/25/99

Jc571 U.S. PTO

Please type a plus sign (+) inside this box → ☐

Approved for use through 09/30/2000. OMB 0651-0032
 Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
 Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PTO/SB/05 (4/98)

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))

Attorney Docket No. BW-02
 First Inventor or Application Identifier Arthur Allen
 Title Method For Connection Acceptance Control
 Express Mail Label No. EJQ 30082079 US

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. ☒ * Fee Transmittal Form (e.g., PTO/SB/17)
 (Submit an original and a duplicate for fee processing)
2. ☒ Specification [Total Pages 21]
 (preferred arrangement set forth below)
- Descriptive title of the Invention
 - Cross References to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference to Microfiche Appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claim(s)
 - Abstract of the Disclosure
3. ☒ Drawing(s) (35 U.S.C. 113) [Total Sheets 12]
4. Oath or Declaration [Total Pages 1]
- a. ☒ Newly executed (original or copy)
- b. ☐ Copy from a prior application (37 C.F.R. § 1.63(d))
 (for continuation/divisional with Box 16 completed)
- i. ☐ DELETION OF INVENTOR(S)
 Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).

* NOTE FOR ITEMS 1 & 13: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).

ADDRESS TO: Assistant Commissioner for Patents
 Box Patent Application
 Washington, DC 20231

5. ☐ Microfiche Computer Program (Appendix)
6. Nucleotide and/or Amino Acid Sequence Submission
 (if applicable, all necessary)
- a. ☐ Computer Readable Copy
- b. ☐ Paper Copy (identical to computer copy)
- c. ☐ Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

7. ☒ Assignment Papers (cover sheet & document(s))
8. ☐ 37 C.F.R. § 3.73(b) Statement of Power of Attorney
 (when there is an assignee)
9. ☐ English Translation Document (if applicable)
10. ☐ Information Disclosure Statement (IDS)/PTO-1449
☐ Copies of IDS Citations
11. ☐ Preliminary Amendment
12. ☐ Return Receipt Postcard (MPEP 503)
 (Should be specifically itemized)
13. ☒ * Small Entity Statement(s) Status still proper and desired
 (PTO/SB/09-12)
14. ☐ Certified Copy of Priority Document(s)
 (if foreign priority is claimed)
15. ☐ Other: _____

16. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No: _____

Prior application information: Examiner _____ Group / Art Unit: _____

For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

17. CORRESPONDENCE ADDRESS

☐ Customer Number or Bar Code Label

(Insert Customer No. or Attach bar code label here)

or ☒ Correspondence address below

Name	Earl Mincer				
Address	4455 E Camelback Suite E160				
City	Phoenix	State	AZ	Zip Code	85018
Country	USA	Telephone	602 840-8411	Fax	

Name (Print/Type)	EARL MINCER	Registration No. (Attorney/Agent)	36,396
Signature	Earl Mincer	Date	6/25/99

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

Jc571 U.S. PTO 09/344688

66/52/99

STATEMENT CLAIMING SMALL ENTITY STATUS
(37 CFR 1.9(f) & 1.27(c))--SMALL BUSINESS CONCERN

Docket Number (Optional)

BW-02

Applicant, Patentee, or Identifier: Arthur Allen

Application or Patent No.: _____

Filed or Issued: _____

Title: Method for Connection Acceptance Control & Optimal Multimedia Content Delivery Over Networks

I hereby state that I am

☐ the owner of the small business concern identified below:

☒ an official of the small business concern empowered to act on behalf of the concern identified below:

NAME OF SMALL BUSINESS CONCERN Instant Video Technologies, Inc.

ADDRESS OF SMALL BUSINESS CONCERN 500 Sansome St Suite 503
San Francisco CA 94111

I hereby state that the above identified small business concern qualifies as a small business concern as defined in 13 CFR Part 121 for purposes of paying reduced fees to the United States Patent and Trademark Office. Questions related to size standards for a small business concern may be directed to: Small Business Administration, Size Standards Staff, 409 Third Street, SW, Washington, DC 20416.

I hereby state that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention described in:

- ☒ the specification filed herewith with title as listed above.
☐ the application identified above.
☐ the patent identified above.

If the rights held by the above identified small business concern are not exclusive, each individual, concern, or organization having rights in the invention must file separate statements as to their status as small entities, and no rights to the invention are held by any person, other than the inventor, who would not qualify as an independent inventor under 37 CFR 1.9(c) if that person made the invention, or by any concern which would not qualify as a small business concern under 37 CFR 1.9(d), or a nonprofit organization under 37 CFR 1.9(e).

- Each person, concern, or organization having any rights in the invention is listed below:
☒ no such person, concern, or organization exists.
☐ each such person, concern, or organization is listed below.

Separate statements are required from each named person, concern or organization having rights to the invention stating their status as small entities. (37 CFR 1.27)

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 CFR 1.28(b))

NAME OF PERSON SIGNING Edward H. Davis

TITLE OF PERSON IF OTHER THAN OWNER Vice President & General Counsel

ADDRESS OF PERSON SIGNING 3616-20th Street SAN Francisco, CA 94110

SIGNATURE Edward H. Davis

DATE 6/24/99

Method for Connection Acceptance Control and Optimal Multi-Media Content Delivery over Networks

Cross-reference to related applications.

This invention claims the priority date of provisional application number 60/108,777, "Method for Connection Acceptance Control and Optimal Multimedia Content Delivery Over Networks", inventor Arthur Allen, filed 11/17/98.

1. Background of the Invention

1.1 Field of the Invention

This invention relates to the field of delivery of multimedia content over a variety of networks. More specifically, it pertains to multimedia servers which service many clients simultaneously for the delivery of multimedia content which is used and played back at each client. It addresses methods for determining optimal delivery rates to each client and methods for determining whether new clients may be accepted without diminishing the quality of service to existing clients.

1.2 Background Information

In the history of multimedia program delivery, some in the industry have long advocated the use of large client-side buffers and faster-than-real-time content delivery over a network as offering the best of all worlds: a jitter-free viewing experience and a cost-effective utilization of the network resources at hand. Few systems, however, go very far in addressing how to schedule clients or a method for accepting new clients. Real-time systems, often known as streaming systems, can schedule new clients in a very simple manner -- if sufficient bandwidth remains for the added real-time stream, then the client may be accepted. However, such systems do not maximize the number of simultaneous clients. On the other hand, faster than real-time delivery, sometimes known as store-and-forward systems, open up the possibility for more flexible scheduling procedures to control and optimize the number of simultaneous clients while ensuring a high level of quality of service.

The methods for such call acceptance and flow modulation that have been proposed in the prior art have been largely ad-hoc and also incomplete. These have been ad-hoc in the sense that there has been no guiding rationale for their selection from among many possible and potentially superior alternatives. The methods have also been incomplete insofar as they did not address the question of whether any given incoming request for service should be accepted or denied. Video-on-demand systems, or more generally, any system in which a multimedia server is designed to serve multiple clients over a network to deliver bounded content, can benefit from the use of such flow modulation techniques and call acceptance procedures.

1.3 Optimal Content Flow Modulation

One time-honored way of designing methods of the class required here is to re-cast the problem to be solved as an optimization problem, in which one seeks to maximize a designated *value function* moment-by-moment, subject to a set of real-world *operational constraints* which will typically vary over time. Accordingly, given a set of clients and associated sessions, an *optimal delivery procedure* continuously establishes content flow rates from the content server to each of its clients so as to maximize aggregate value according to the governing value function.

This approach holds several advantages: 1) optimization problems are well understood, and are tractable by a large and diverse collection of computational methods; 2) if it exists, the global solution that is obtained is arguably optimal by construction, and thus superior or equal to all other.

The present invention teaches the method of optimizing two particular value functions:

- 1) total data delivered. (maximize throughput)
- 2) total delivery charges (maximize charges).

The first value function does not distinguish one customer from another and will deliver as much data as possible from server to clients irrespective of the characteristics of the latter. The second value function favors the service of high paying customers. It can easily be seen that the first function is a special case of the second one whereby all clients are charged equally.

As will be seen in this disclosure, optimizing for these functions and identifying the necessary constraints requires a new and unique perspective that is specifically designed for the multimedia environment. Moreover, the disclosed methods are specifically designed to account for and accommodate real-world scenarios of today's networks. Consequently many variations of the method are presented to accommodate various scenarios.

2. Summary of the Invention

2.1 Call/Connection Acceptance Control (CAC)

A CAC procedure is responsible for deciding whether a candidate for service can be accommodated without jeopardizing sessions already in progress at the present time or at some time in the future; failing that it must decide whether a service request should be queued for a time, or rejected.

2.1 Flow modulation

Flow modulation methods are those portions of the system which manage the communication and data flow between the server and the clients. Collectively, these methods provide the multimedia data to the client and provide the server with the information about the state of the transmission, playback, user status and network status. These parameters are further used by the present invention in the CAC procedures. In fact, as will be shown, the proposed CAC procedures are tightly integrated with the flow modulation methods.

2.3 Adaptation to Variations in Network Capacity

Operational constraints may change over time. For instance, one might elect to vary the total bandwidth available for multimedia content delivery according to the time of day. Alternatively, exogenous data flows on the network may cause unexpected disturbances by

usurping available bandwidth and impeding the delivery of data along established session channels. The content delivery strategy of the present invention includes the ability to adapt to scheduled as well as unexpected disturbances so as to minimize unwanted disruptions of services.

2.4 Burst transmissions provide the opportunity to adapt

The present invention, due to its faster-than-realtime transmissions (also known as burst transmissions), which are realized by use of high-bandwidth networks and large client cache or intermediate storage, provides an opportunity to adapt to changing network conditions. In contrast real-time (streaming) systems are essentially designed for worst-case scenarios: each client must be assumed to constantly use the complete real-time playback bandwidth. Such a system is unable to adapt to any derivation from this scenario. For example, take the simple case where the total server bandwidth is 100% utilized by all clients playing back the streaming video. Should any network condition change, such as a temporary decrease in available bandwidth over the network, then one or more clients' playback is interrupted, and the system can not recover from such a condition until the bandwidth is regained. Even worse if a single client presses pause either that unused bandwidth must remain reserved and no more clients can be accepted, or that paused client is pushed out in order to service the new client. In essence little or no CAC procedure may be implemented.

In contrast the present invention burst transmits portions of a program and immediately 'gets ahead of itself', thus allowing headroom for a myriad of methods to intelligently handle new clients, client interactivity and possible network fluctuations.

Methods are taught for optimally determining the flow rate to each client. Methods are also taught for accepting or rejecting new clients; these call-acceptance methods are tightly coupled with said flow rate modulation methods. A series of constraint expressions are presented which govern the methods for determining the flow rates and acceptance of new clients. Linear programming techniques are used to optimally solve these expressions. Various embodiments are presented including scenarios for multiple-rate tariffs, and time-of-day bandwidth variations.

3. Brief Description of the Drawings

Figure 1 depicts the flow of control and/or data between the different stations of a content delivery session.

Figure 2 illustrates the Entity Data Model.

Figure 3 geometrically illustrates the problem statement.

Figure 4 geometrically illustrates an expansion of the problem statement.

Figure 5 illustrates a method for implementing flow modulation.

Figure 6 illustrates a method for implementing flow modulation for maximized charges.

Figure 7 illustrates typical content flow.

Figure 8 illustrates typical server swing capacity.

Figure 9 illustrates a method for call-acceptance and control (CAC).

Figure 10 illustrates planned constraints on maximum flow.

Figure 11 illustrates a method for call-acceptance and control (CAC) with scheduled flow changes

Figure 12 illustrates stratification of services.

Figure 13 illustrates a method for call-acceptance and control (CAC) for maximal charge.

4. Detailed Description of the Invention

4.1 Data & Control Flows (Figure 1)

Figure 1 depicts the flow of control and/or data between the different stations of a content delivery session. As shown a client attempts a connection (100) and manifests itself to the Content Selection subsystem by means of a low bandwidth control channel (not shown). Next the client is authenticated and a selection is made (110), typically with the aid of a browser software such as Netscape or Microsoft Internet Explorer. If the client is not authenticated, it is dismissed from the system (120). If the client has been authenticated and a program selected for viewing then the rate of service is set at this point (130), perhaps according to the selection that was made, or some contractual stipulation. The client is now placed on the service queue of the CAC subsystem (140). A client that is made to wait too long will eventually balk (150). Assuming this does not occur, the CAC subsystem will eventually allocate a channel to the client and open a session (160). Control now devolves upon the Content Flow Modulator (not shown) which starts the flow of content from server to client (170). Subsequent capacity changes, whether predictable or not, may force an abrupt termination of a session in progress (180). Otherwise the session runs to completion (190).

4.2 Entity Data Model (Figure 2, listing, table)

The entities entering into our discussion are depicted in figure 2. Client 200 maintains certain data associated with this entity; as shown but not labeled, which includes without limitation, status, id and costOfService. The other entities also each include unlabeled but depicted data. The diagram further depicts the relationship between each entity. As shown, client 200 is assigned a session 240. Client 200 employs a channel 210. Client 200 selects contentSelection 230. Session 240 delivers content through channel 210. Server 220 modulates channel 210. Server 200 contains contentSelection 210. Server 220 accepts, defers or denies client 200. And contentSelection 230 is associated with session 240.

Furthermore the figure depicts the various one-to-many relationships. Each client 200 employs one channel 210. Client 200 may or may not receive one of channel 210, as depicted by the 0/1 notation. Similarly, client 200 may or may not receive a session 240. However, whenever client 200 does receive a session 240, it will always receive a channel 210 since channel 210 and session 240 are allocated as a pair. One or more (N) of client 200 may select one of contentSelection 230. And server 220 contains one or more (N) of contentSelection 230. Each one of contentSelection 230 is associated with 0-N of session 240. Each session 240 delivers content through one of channel 210. And server 220 modulates one or more (N) of channel 210.

A more detailed list of each entity of figure 2, and each one's associated description, data elements and function calls is listed below. This listing closely resembles that of object-oriented programming. As such, 'methods' represent the ability to obtain or modify data, while 'attributes' represent data which is directly associated with that particular entity. The listing also includes information relating to one embodiment wherein software programming specifics are disclosed, such as a variable type (double, int and so forth) and more. The present invention is not limited to such an embodiment and other implementations are possible without deviating from the scope and intent of the present invention. The listing, however detailed, is merely illustrative of the data and functions which are used in the equations and methods described herein.

Consequently, data and functions from this listing, associated with the various entities, will be used in forthcoming equations, flowcharts and methods. The reader is directed to this listing as reference when reading such equations and examining such drawings.

--- start of entity data model detailed listing ---

Model: Untitled 1 (public)

Contains:

client, session, channel, server, contentSelection.

Component: client (public Class/Interface)

Comment:

A client entity stands for a client presently requesting or receiving service.

Methods:

```
public static lookup (id: in int) : client
public GetId () : const int&
public SetId (val : in int&)
public GetCostOfService () : const double&
public SetCostOfService (val : in double&)
```

Attributes:

```
private status: client<int>
Specifies whether or not a client has been allocated a channel and session.

private id: int
Integer-valued identifier that is unique to the client (primary key). Can be obtained from a monotonically increasing counter.

private costOfService: double
Dollar charge per Mbyte. This value is the same for all customers under flow optimization. Under cost/charge optimization may be an integer value reflective of the rank; the higher the rank the higher the charge.
```

Has:

```
public selected: contentSelection
public assigned a: session
public employs: channel
```

Component: session (public Class/Interface)

Comment:

A session entity holds various state information about the service being received by an associated customer.

```
public GetCurrentPosition () : const double&
public SetCurrentPosition (val : in double&)
public GetPayloadToGo () : const double&
public SetPayloadToGo (val : in double&)
public GetStatus () : const int&
public SetStatus (val : in int&)
public GetMinimumFlowRate () : const double&
public SetMinimumFlowRate (val : in double&)
public GetFlowRateRange () : const double&
public SetFlowRateRange (val : in double&)
public GetMaxFlowRate () : const double&
public SetMaxFlowRate (val : in double&)
```

Attributes:

```
private playTimeToGo: double
```

Indicates the minutes remaining in the viewing experience. Initialized to `contentSelection.playTime` (see below).

private currentPosition: double

Pointer into media content from which content is being delivered.

private payloadToGo: double

The amount of media content (in Mbytes) as yet undelivered by the server. Does not include any content presently stored in the client-side buffer..

```
private status: int
```

Indicates whether session is active or paused.

```
private minimumFlowRate: double
```

This is the minimum flow from server to client required to ensure uninterrupted service over the remaining playTime. Has a value of zero if payloadToGo is zero. Given by $(\text{payloadToGo} * 8) / (\text{playTimeToGo} * 60)$

```
private flowRateRange: double
```

Specifies the effective range over which the channel content flow serving a session is constrained without consideration for interactions with other flows. Equals $\text{maxFlowRate} - \text{minimumFlowRate}$

```
private maxFlowRate: double
```

Effective maximum bound on flow as expressed in formula (8) which must be re-evaluated periodically.

Has:

public delivers content through: channel

Component: channel (public Class/Interface)

Comment:

A channel represents the network resources from server to client associated with an ongoing session, encompassing the client-side buffer if any, and its level.

public GetBufferLevel () : const double&**public SetBufferLevel (val : in double&)**

```
public GetFlowRate () : const double&
```

```
public SetFlowRate (val : in double&)
```

```
public GetMaxFlowRate () : const double&
```


public SetMaxFlowRate (val : in double&)

Attributes:

private bufferSize: double

Capacity of the client-side buffer (or equivalent)

private bufferLevel: double

Current buffer level in MBytes of stored content.

private flowRate: double

Flow rate through channel specified by the relevant optimizing flow modulator.

private maxFlowRate: double

This value represents the maximum possible flow rate from the server to an individual client over its "channel". This value reflects restrictions on flow that pertain to an individual client. It may be determined by factors such as the bandwidth of client's link to the network, or a limit imposed administratively to ensure balanced network utilization..

Component: server (public Class/Interface)

Comment:

Entity representing the media server and its CAC and flow modulation activities.

public GetFlowRate () : const double&

public SetFlowRate (val : in double&)

public GetMaxMinFlowRate[] () : const double&

public SetMaxMinFlowRate[] (val : in double&)

Attributes:

private maxFlowRate: double

Maximum possible content flow that is allocated to the server by the network.

private flowRate: double

Aggregate content flow rate, summed over all sessions and their associated channels.

private cac_flowSafetyMargin: double

Tunable safety margin used by the CAC algorithm to protect sessions-in-progress from being affected by changes in available network bandwidth.

private maxMinFlowRate[]: double

Applies when N rate tariffs exist. This array holds the maximum floor level for each category of service. The value for the costliest category N is stored in maxMinFlowRate[N-1], and for the least costliest in maxMinFlowRate[0]. It is the relative magnitude of these ascending values that matters, not their absolute value. Thus the actual maximum floor flow rate for category k is given by $\text{server.maxFlowRate} * (\text{server.maxMinFlowRate}[k-1] / \text{server.maxMinFlowRate}[N-1])$. Similarly, the maximum floor flow rate for category N is $\text{server.maxFlowRate}$.

Has:

public contains: contentSelection

public modulates: channel

Component: contentSelection (public Class/Interface)

Comment:

Entity represents a video/sound clip or other bounded unit of content. A continuous data feed does not qualify.

Attributes:

private averagePlayRate: double

[illegible]

total size of the content in Mbytes.

public is associated with: session

The following table summarizes the highlights of the previous detailed description of each entity in Figure 2.

page 8 of 37

4.3 Constraints On Content Flow

Before moving on to more figures, it is imperative to establish some formulas and problem statements which are used in the methods which follow.

The flow of content between entities is subject to the following constraints at all times. Buffer levels are always expressed in Mbytes and data rates in Mbits/sec.

$$(1) \sum_{i \in S_{\text{activeClients}}} (\text{client.lookup}(i).\text{channel.flowRate}) \leq \text{server.maxFlowRate}$$

The sum of all channel flows cannot exceed the imposed maximum throughput capacity of the server.

$$(2) \text{client.lookup}(i).\text{channel.flowRate} \leq \text{client.lookup}(i).\text{channel.maxFlowRate} \\ \text{for all } i \in S_{\text{activeClients}}$$

The data path from server to client is subject to its own constriction.

$$(3) \text{client.lookup}(i).\text{channel.flowRate} \leq \\ (\text{client.lookup}(i).\text{channel.bufferSize} - \text{client.lookup}(i).\text{channel.bufferLevel}) * 8/60 \\ + \\ \text{client.lookup}(i).\text{session.mediaContent.averagePlayRate}, \\ \text{for all } i \in S_{\text{activeClients}},$$

The channel buffer is never allowed to overflow.

$$(4) \text{client.lookup}(i).\text{channel.flowRate} \leq \text{client.lookup}(i).\text{session.payloadToGo} * 8/60 \\ \text{for all } i \in S_{\text{activeClients}},$$

Content that does not exist can not be delivered. (Constraint 1 will ordinarily prevail except at the very end of a session.)

The constraints listed above are straightforward applications of common sense in relation to the flow of data through constricted channels, out of finite data sources, and into and out of bounded buffers. By contrast, the following constraint, which imposes a minimum channel flow rate instead of a maximum, is less obvious. Indeed its introduction here and subsequent use in the CAC procedure constitute the central innovation of the methods presented in this document. The minimum value, termed the *minFlowRate* is set to the flow rate which, if sustained over the balance of the play time to go (*playTimeToGo*), ensures that all required content will be available when needed -- and no sooner -- until all content is played out. This floor value can be calculated for $i \in S_{\text{activeClients}}$ by the formula

$$(5) \text{client.lookup}(i).\text{session.minFlowRate} = (\text{client.lookup}(i).\text{session.payloadToGo} * 8) / \\ (\text{client.lookup}(i).\text{session.playTimeToGo} * 60)$$

Thus:

$$(6) \text{client.lookup}(i).\text{channel.flowRate} \geq \text{client.lookup}(i).\text{session.minFlowRate} \\ \text{for all } i \in S_{\text{activeClients}}$$

The variable constraint bounds (i.e. the values to the right of the inequality symbol) of equations 1 – 4 and 6 are re-evaluated on a periodic basis (e.g. once per second) prior to the

execution of the CAC procedure and optimizer. In particular, the *minFlowRate* value starts out at the beginning of a session equal to the streaming rate. By construction the *minFlowRate* rate never exceeds this initial value so long as constraint 6 is honored. In fact, constraint 5 implies that the *minFlowRate* value must be a diminishing function of time that may hold its value for a time but never rises. As seen from equation 6, the actual data rate of the channel, *flowRate*, is always greater than or equal to the *minFlowRate*. By design, and virtue of the fact the present invention uses faster-than-realtime transmissions, the system quickly gets ahead of itself and ensures that after initial conditions, the *minFlowRate* is always equal to or less than the real-time rate and that it continues to decrease. As we shall see the CAC procedure exploits this monotonic characteristic of the minimum flow rate over time.

Constraints 2, 3 and 4 are of like kind, each specifying an upper bound on individual channel flows. Whereas the bound for constraint 2 is typically a constant, the bounds on 3 and 4 will vary over time. Nevertheless, only one of the three bounds is effective at any given time, namely the one with the smallest bound value, given by:

(7) $\text{client.lookup}(i).\text{session.maxFlowRate} = \text{minimum of}$

- 1) $\text{client.lookup}(i).\text{channel.maxFlowRate},$
- 2) $(\text{client.lookup}(i).\text{channel.bufferSize} - \text{client.lookup}(i).\text{channel.bufferLevel}) * 8/60 + \text{client.lookup}(i).\text{session.mediaContent.averagePlayRate},$
- 3) $\text{client.lookup}(i).\text{session.payloadToGo} * 8/60$

Consequently, formulas 2, 3, and 4 can be consolidated into a single constraint, the bound for which is computed at every scan to be the smallest bound of associated constraints 2, 3 and 4.

(8) $\text{client.lookup}(i).\text{channel.flowRate} \leq \text{client.lookup}(i).\text{session.maxFlowRate},$

whereby for all $i \in S_{\text{activeClients}}$, *maxflowRate* is given by equation (7).

At any one time, individual channel flows are constrained over a range, as follows:

(9) $\text{client.lookup}(i).\text{session.flowRateRange} = \text{client.lookup}(i).\text{session.maxFlowRate} - \text{client.lookup}(i).\text{session.minimumFlowRate}$

4.3.1 Value Functions

The value functions introduced in section 1.2 can now be expressed mathematically as linear functions of channel flows, as follows:

4.3.1 Optimizing Throughput (Maximal Flow)

(10) $\text{value} = \sum_{i \in \text{ActiveClients}} \text{client.lookup}(i).\text{channel.flowRate}$

4.3.2 Optimizing Charges (Maximal Charges)

(11) $\text{value} = \sum_{i \in \text{ActiveClients}} (\text{client.lookup}(i).\text{channel.flowRate} * \text{client.lookup}(i).\text{costOfService})$

4.3.3 Optimization Problem Statement (Figure 3)

The optimization problem, which in one embodiment is strictly throughput and in another case is charge, can be stated simply as follows:

Find values for

$\text{client.lookup}(i).\text{channel.flowRate}$ for all $i \in S_{\text{activeClients}}$

constrained by inequalities 1 through 5, such that the value obtained by evaluating expression 10 or 11 assumes a maximum.

Both of these problem formulations are examples of Linear Programming for which a number of well-known and generally effective computational solutions exist. In linear programming one seeks to optimize a linear cost function of variable x

$$(12) \quad c^*x = c_1^*x_1 + \dots + c_n^*x_n$$

subject to a set of linear inequality constraints

$$(13) \quad A^*x \leq b$$

where $x^T = (x_1, \dots, x_n)$, $c = (c_1, \dots, c_n)$ are the state variable & cost vectors, A is an n -by- m matrix, $b^T = (b_1, \dots, b_m)$ is the constraint vector, and the operator '*' stands for matrix or scalar multiplication.

Figure 3 is introduced as illustrative of the problem statement and the general methods of the prior art, and is not incorporated as an element of the invention.

The linear programming problem as well as its solution can best be understood with the aid of geometry. Figure 3, depicting a 2-dimensional Cartesian problem space, inequality constraints (13) define a convex hull H 310 over which a search for an optimum value of $x = (x_1, x_2)$ is permitted to range. The cost vector c 350 defines an infinite family of equal cost lines (hyper-planes) which lie orthogonal to c . Three examples of such lines are shown in L_1 360, L_2 365, and L_3 370, each of progressively higher value. The supreme value of the cost function is obtained by sliding along c 350 till one can go no further, in this instance toward vertex V_4 340 of hull H 310. Many well-known methods (e.g. the Simplex Method) work roughly in this fashion, exploiting the fact that at least one optimum point must be at a vertex. In particular, the Simplex method algorithm begins by finding a vertex (e.g. V_2 320), and then moves along a sequence of vertices (e.g. V_3 330, V_4 340) improving itself each time until no further improvement is possible & the summit is reached.

Let us suppose instead that V_3 330 were placed along L_3 370 along with V_4 340. According to prior art methods, V_3 330 and V_4 340 are the two possible solutions, but the equally valuable points in between them are not. As we shall soon see, the problem of throughput optimization (6) falls in this category.

While vertex V_1 300 does not factor into this description, it is depicted in figure 3 for completeness.

4.4 Flow Modulation

4.4.1 A Method for Maximal Flow

The following sections detail one embodiment to meet object 1 of section 1.3, namely to optimize the total data flow.

4.4.1.1 Overview (Figure 4)

Figure 4 depicts a scenario involving two flows. The convex hull is in this instance bounded by line segments L1, L2, L3, L4 and L5. L6 is a boundary used in a different embodiment, however the present embodiment uses L5 and not L6. Flow f_2 can range over the interval separating line segments L1 from L3, namely f_2^{MIN} and f_2^{MAX} , the range is depicted as f_2^{RANGE} . Flow f_1 can range over the interval between lines L2 and L4, namely f_1^{MIN} and f_1^{MAX} , and depicted as f_1^{RANGE} . The sum of flows f_1 and f_2 is constrained to lie inside of line segment L5 which, by construction, is always orthogonal to the cost vector C_f . Cost vector C_c is also illustrated but used in a distinct embodiment. In the present embodiment, only C_f is used. In the illustrated example of the present embodiment the constraint on total flow is set to 5, and is therefore low enough to cause L5 to intersect L3 and L4. This would not have been true had the value chosen had been 10 instead of 5. With L5 safely out of contention, the convex hull would instead be a simple rectangle bounded by L1 through L4, thereby permitting both flows to assume their respective maxima without interference. In practice operational constraints exist intrinsically or are imposed from the outside so as to ensure cost effective sharing of potentially costly network resources.

Supposing figure 4 to be correct, the well-known methods would select would vertex V_{3-5} , which lies at the intersection of L3 and L5, or V_{4-5} , which lies at the intersection of L4 and L5. These solutions, though optimal, are undesirable for the present invention as they fail to spread available bandwidth over all channels as fairly as would a centrally located interior point of L5. For this reason a simple optimization method is taught, which is adapted to the particular needs of this problem and ensures a fairer allocation of constrained bandwidth among all channels.

4.4.1.2 Procedure (Figure 5)

In order to optimize use of all available bandwidth, the following general method is used, with the details illustrated in figure 5 This method is a solution for the problem illustrated in Figure 4, which geometrically illustrates the optimization problem in the limited case of two flows, f_1 and f_2 . The following description expands the problem to an arbitrary number of clients (and therefore flows) and presents a method for solving this optimization problem.

Referring to figure 5, in step 500 values are calculated for the *session.maxFlowRate* and *session.minFlowRate* for each client as per the minimum and maximum constraint bound expressions in 6 and 8, respectively.

The difference between these two yields the *session.flowRateRange* of each client. Thus

$$\text{session.flowRateRange} = \text{session.maxFlowRate} - \text{session.minimumFlowRate}$$

In step 505, the active clients are sorted in an ascending fashion based upon their *session.flowRateRange*. As will be shown this critical step facilitates allocation of the remaining server bandwidth as evenly as possible among all active channels, thus maximizing the number of channels that benefit by use of the total server bandwidth. An arbitrary assignation of remaining bandwidth is likely to saturate the server before all channels have been assigned extra bandwidth, thereby favoring certain channels on an ad-hoc basis.

In step 510, each client's channel flow rate is set to the *session.minimumFlowRate*.

By doing so it is ensured that the minimum flow constraint is met for each session and that the minimum flow rate is a non-increasing function of time, which is critical to the proper functioning of the CAC procedure. All clients are marked as unprocessed.

In the next step, 520, *server.flowRate* is set to the sum of each active client's *session.flowRate*.

Next, the following is iterated over all clients in sorted sequence (during any given iteration the selected client is given by its *id*) by performing steps 530 through 570. In step 530 evaluating the following expressions test for possible server saturation:

$$\begin{aligned}\text{delta} &= (\text{server.maxFlowRate} - \text{server.flowRate}) / (\text{qty of un-processed clients}) \\ \text{range} &= \text{client.lookup}(\text{id}).\text{session.maxFlowRate} - \text{client.lookup}(\text{id}).\text{session.flowRate}\end{aligned}$$

If *range* is greater then *delta*, this implies that the server can be saturated in this iteration by allocating *delta* to all unprocessed clients (step 540).

On the other hand, the 'no' path for step 530 implies that the server is not saturated and that the present client (given by *id*) will saturate first. Accordingly, in 550 the *delta* variable is set as follows:

$$\text{delta} = \text{range}$$

Next, the flow rate is incremented for all unprocessed clients by *delta*, causing client *id* to saturate.

In step 560 the server flow rate is adjusted accordingly:

$$\text{server.flowRate} = \text{server.flowRate} + \text{delta} * (\text{qty of unprocessed clients})$$

In step 570 the client given by *id*, now saturated, is marked as processed.

4.4.2 A Method for Maximal Charge

The following sections detail one embodiment to meet object 2 of section 1.2, namely to optimize the total monetary charges within the system.

4.4.2.1 Overview (Figure 4)

Referring back to Figure 4, cost vector C_c lies orthogonal to line L6, which intersects the convex hull at the vertex formed by the intersection of lines L4 and L5, namely $V_{4.5}$. This cost vector, and the optimal point that it implies, favors flow f1 over flow f2. In this example, this is as it should be, as the cost of service for f1 equals 2, thus exceeding the cost of service of 1 set for f2. As the number of flows grows to exceed the number of distinct categories of service (and associated costs of service) the unique optimal solution, depicted in figure 4 for the case where every flow has a distinct cost of service, no longer applies. Once again a plurality of flows within a service category vie for bandwidth which a method should endeavor to distribute evenly. This method is derived from the previous one, and optimizes one cost category after another, starting with the most costly and ending with the least costly, or when all available capacity is allocated.

4.4.2.2 Procedure (Figure 6)

Let the service categories be denoted by $k=1..N$, where k also denotes the cost of service.

Let $C_1..C_N$ be the partition of $S_{\text{activeClients}}$ that places all clients with cost of service k in set C_k . Partition sets C_k can be ordered to form sequence $\text{SeqC} = C_N..C_1$.

Figure 6 depicts the method for implementing the method to maximize the cost of service (service charge) according to objective function 2 in section 1.3.

This method is nearly identical to the previous one. The principle difference stems from the partitioning of clients according to their category (cost) of service: clients charged most are allocated bandwidth preferentially. This is accomplished by adding another level of iteration around the method of figure 5. The inner iteration (steps 650 through 680) functions exactly as before, with the difference that its actions are limited to the clients belonging to the given service category k (i.e. C_k). This difference also holds true of step 640 which sorts category k clients according to their flow ranges prior to entry in the bandwidth-allocating inner loop. The outer loop proceeds down a sorted sequence of service categories SeqC (generated in step 630), starting with the category generating the greatest revenue to the service provider. Given a fairly static set of service categories, this sort need be performed only when the categories undergo change. Steps 670, 675 and 680 are identical to their counterparts in the method of figure 5 (i.e. 570, 575 and 580).

The net effect of this method is preferential allocation of bandwidth according to category of service, and equitable treatment of clients within the same category of service.

4.5 Call Acceptance Control (CAC)

4.5.1 CAC for Maximal Flow

4.5.1.1 Overview (Figures 7-8)

The CAC procedure applicable to this flow optimization relies on the essential step of accepting a new client if and only if the added load induced thereby does not compromise service to existing clients or the new one. This critical step could not be accomplished without the close integration with previously- described flow-modulation methods of figures five and six.

According to the previous discussion, the minimum flow rate is the minimum sustained flow rate that guarantees that the associated viewer will not be subject to interruptions in service due to a shortfall of content from the media server. It follows that whenever data is being delivered at a rate in excess of the minimum flow rate, a downward adjustment toward the minimum level could be accommodated as needed to surrender bandwidth to any newcomer.

Figure 7 depicts content flow over a channel over the course of a typical session, and also how data is delivered under real-time streaming (D). The amount of content delivered is the same in either case, but the manner of delivery differs considerably. A session is launched at time 0 as the network is lightly loaded, and the optimizer sets an accordingly high flow rate. Another client emerges at the end of interval 700, causing a downward adjustment to the flow rate over interval B, as available bandwidth is shared between two sessions. During both of these intervals the minimum flow rate 720 drops quickly, as data accumulates in the client's media buffers. At the end of interval B a massive influx of clients necessitates that flow be dropped to the minimum flow rate, which now lies substantially below the streaming rate D and is held until all data is delivered at the end of interval C. Note that the minimum flow rate, graphed as element 720, diminishes monotonically over time.

The *server swing capacity* is defined as the difference between the maximum capacity of the server and the total minimum flow rates for all active clients. Therefore:

(14) swingCapacity =

$$\text{server.maxFlowRate} - \sum_{i \in \text{SActiveClients}} (\text{client.lookup}(i).\text{session.minFlowRate})$$

Given the monotonic decreasing nature of session minimum flow rates, server swing capacity can readily be seen to be a *monotonic increasing* function of time over the intervals separating client admissions, at which points it undergoes a drop as a new load is taken on. This all-important characteristic implies the following:

Any client admitted for service based on the present value of swing capacity is guaranteed to have sufficient bandwidth at its disposal over the entire future course of the session.

Figure 8 depicts the server swing capacity 800 over the course of the sessions illustrated in figure 7. Swing capacity rises quickly over intervals A & B as data is delivered at high flow rates over the network. It holds steady over interval C when all channels flow at their minimum rate then jumps at the end of C before resuming its monotonic rise once again.

4.5.1.2 Procedure (Figure 9)

In this procedure, which executes on a periodic basis, queued clients awaiting bandwidth are scanned in FIFO order. For each one the required bandwidth is computed as per the client's prior content selection. If the available swing capacity reduced by a safety margin exceeds the amount required then the client is activated, and swing capacity is adjusted accordingly. Otherwise two possible cases are considered: 1) under the *FirstFit* embodiment the procedure continues scanning clients to the end of the queue, activating clients whose requirements can be met; 2) under the *FIFO* embodiment, the procedure ends with the first candidate client whose requirements cannot be met.

In step 900 available server swing capacity is evaluated according to the formula

$$\text{swingCapacity} = \text{server.maxFlowRate} - \sum_{i \in \text{SActiveClients}} (\text{client.get}(i).\text{session.minimumFlowRate})$$

The bandwidth requirement for client *id* in Step 920 is obtained as follows:

$$\text{required_bandwidth} = \text{client.lookup}(\text{id}).\text{contentSelection.averagePlayRate}$$

The predicate evaluated in Step 940 is given by the expression

$$(\text{required_bandwidth} \leq \text{swingCapacity} - \text{server.cac_flowSafetyMargin})$$

In step 950, client activation entails allocation of a session and a channel, and insertion in the set of active clients eligible for bandwidth allocation by the optimal flow modulator.

In step 960 the swing capacity is diminished by the amount reserved for the activated client:

$$\text{swingCapacity} = \text{swingCapacity} - \text{required_bandwidth};$$

4.5.2 Responding to Variations in Network Capacity (Maximal Flow)

In the CAC procedure for maximal flow (section 3.5.1), a safety margin was introduced, namely *server.cac_flowSafetyMargin*, to provide the means for ensuring that the server's swing capacity will never fall below a minimal threshold value.

According to this procedure, the following inequality always holds true:

$$(15) \quad \text{swingCapacity} \geq \text{server.cac_flowSafetyMargin}$$

In the previous discussion a server's swing capacity provided the basis for determining whether or not a prospective client should be allocated bandwidth. In another embodiment, server swing capacity can also be interpreted as specifying the *maximum* amount by which the *server.maxFlowRate* constraint can be dropped without affecting service, should such an adjustment prove necessary due, for instance, to an influx of exogenous network traffic that diminishes the amount available for multi-media services. Parameter *server.cac_flowSafetyMargin* can thus be set so as to *guarantee a minimum capacity to tighten the constraint on maximum server flow* in response to unexpected load changes that affect the server's ability to service its existing clients as well as new ones.

4.5.3 Anticipating Scheduled Variations in Network Capacity (Maximal Flow)

4.5.3.1 Overview (Figure 10)

Figure 10 depicts how the constraint on maximum flow might be allowed to vary according to the time of day, day of the week, and so forth, in expectation of time-varying traffic flow levels extrapolated from past experience, traffic flow models, etc. Maximum flow rate 1000 can be seen to vary based upon the time of day. In practice, defining the right-hand-side of inequality constraint 1 as a time-dependent expression can impose such time-varying capacities. According to the previous description, the optimizer, which executes on a periodic basis, will automatically seek new flow levels for every active session as the constraint varies. There is, however, no guarantee that an acceptable operating point will be found for all sessions (i.e. one that respects the minimal and maximum constraints on session channel flow). One such example is the case where the server is loaded to the limit and total capacity is curtailed in excess of the aforementioned safety margin. Should such a situation arise the only recourse may well be the termination of a number of established sessions (i.e. load shedding).

The goal is to eliminate service disruptions of this sort by allowing the CAC procedure to look ahead into the future, and accept new clients only if these can be accommodated without any compromise in service in the midst of *previously anticipated changes* in available network capacity. The following CAC procedure generalizes the previous one: before accepting a client the test on swing capacity is repeated over a sequence of time segments that cover the proposed viewing period.

4.5.3.2 Definitions

Let

$$(16) \quad t_{\text{end}}(i) = \text{client.lookup}(i).\text{session.playTimeToGo} + t_{\text{now}}$$

Let *server.maxFlowRate(t)* be server flow capacity as a function of time, as exemplified in figure 10.

Let *Seq_T(t_{now})* = advancing sequence of future times, lead by *t_{now}*, when *server.maxFlowRate(t)* undergoes a step change. For instance, at 9:15 in figure 10 this sequence reads as follows: 9:15, 9:30, 11:30, 13:30, 6:30, 7:30.

The server swing capacity at a future time *t* is computed according to the capacity and worst-case client flows at time *t*.

Note that only active clients whose t_{end} times occur after t are considered in the sum of minimum flow rates.

The predicate expression used in step 1140 at time t is thus

$$(\text{required_bandwidth} \leq \text{swingCapacity}(t) - \text{server.cac_flowSafetyMargin})$$

Step 1160 performs the same actions as step 660 in the previous cac flowchart

The first CAC process presented in section 4.5.1 is a special case of the present one, in which the set of step change times to $\text{server.maxFlowRate}$ is empty (i.e. $\text{server.maxFlowRate}$ is constant), and $\text{Seq}_T(t_{now}) = t_{now}$.

In preparing $\text{Seq}_T(t_{now})$, one need only consider future times that will pass before the longest possible content is played out if started at t_{now} . In order to sidestep problems associated with rollover (at midnight, year 2000, etc.), time is best expressed as a monotonically increasing value (e.g. seconds since Jan 1 1990).

4.5.3 CAC for Maximal Charges

4.5.3.1 Overview (Figure 13)

In section 4.4.2 a method for flow modulation was presented that maximizes charges to clients with active sessions. The CAC embodiments presented previously is not sufficient as it does not consider the cost of service as a basis for connection acceptance. As a result it may turn away higher paying customers while granting service to lower paying ones, thereby defeating the purpose for this particular embodiment. Therefore, in another embodiment is defined which offers the following features:

1. Awaiting clients are serviced in order of their respective service categories, higher paying clients first.
2. Once accepted, a client is guaranteed to receive acceptable service irrespective of its service category.
3. Under heavy load conditions higher paying customers are more likely to be accepted than lower paying ones.
4. Lower paying customers will not be starved for service when higher paying ones enjoy a surplus.
5. Available bandwidth is not thrown away needlessly while clients are being denied service.

The first objective is easily met by dividing the client queue into as many bands as there are service categories, resulting in a banded queue. Bands are ordered within the queue according to their service categories, with the costliest category in front. As prospective clients arrive and make their selection they are placed in their respective queue band according to their service category (which may be set contractually, according to content selection, etc.).

Our second objective is met by employing a procedure patterned after those presented previously & offering the same guarantee. Toward our third and fourth objectives we propose dividing total available bandwidth in as many strata as there are service categories, as depicted in Figure 12 Two service categories are shown, Premium and Basic, each entailing an

associated cost of service. A prospective client is accepted only if there is sufficient swing capacity available within its service category. The swing capacity for a given category is given by the smaller of 1) the difference between its maximum floor flow rate -- corresponding to the top of the stratum for the service category -- and the sum of the minimum rates of all active sessions in its category or below, and 2) available swing capacity overall. Finally, our fifth objective is met by allowing the flow optimizer to function freely subject to its operational constraints. The imposed ceilings on call acceptance by category relate to *minimum flow rates*, which merely impose a floor on *actual flow rates*. For example, basic clients might well consume all available bandwidth (300) in the absence of any premium customers, yet could be throttled back toward their floor flow rates (which together cannot exceed 200 in this example) at any time should any premium customer suddenly demand service. In contrast, premium customers could consume the entire 300 bandwidth. As lower paying customers appear these would be admitted to the degree that their quota on minimum flow is not exceeded (i.e. 200) and the availability of swing capacity on the system.

4.5.3.2 Procedure (Figure 13)

The present procedure requires a number of ancillary definitions, which follow:

Let the service categories be denoted by $k = 1..N$, where k also denotes the cost of service.

Let $\text{server.maxMinFlowRate}[k-1]$ be the top of the stratum for service category k . Note that $\text{server.maxMinFlowRate}[N-1] = \text{server.maxFlowRate}$.

Let S_k be the set of active client indices with a service category *equal to or less than* k . Note that S_1 is contained in S_2 , S_2 is contained in S_3 , and so forth, and that $S_N = S_{\text{ActiveClients}}$.

Let $\text{swingCapacity}(k)$ denote available swing capacity for service category k . By construction,

$$(19) \quad \text{swingCapacity}(k) = \text{minimum of :} \\
\begin{aligned}
& (\text{server.maxMinFlowRate}[k-1] - \sum_{i \in S_k} (\text{client.lookup}(i).\text{session.minFlowRate})), \\
& (\text{server.maxFlowRate} - \sum_{i \in S_{\text{ActiveClients}}} (\text{client.lookup}(i).\text{session.minFlowRate}))
\end{aligned}$$

Now, referring to figure13:

This method is used for CAC when multiple rate tariffs are in effect, and there is a desire to maximize economic returns to the service provider while offering acceptable service to all.

All waiting clients are scanned in FIFO sequence. The actions taken in Steps 1320 and 1360 are identical to those described in connection with earlier CAC flowcharts.

Step 1340 evaluates a predicate expression that tests whether the required bandwidth can be accommodated without exceeding the lesser of 1) swing capacity available to the client's category of service, and 2) total available swing across all categories of service. The latter factor could be determinative if all available bandwidth were allocated to high paying customers, leaving lower paying ones such as the proposed client unable to draw from their unfilled quota.

Let us suppose that candidate client *id* belongs to rate category k .

We define the swing capacity available in rate category k as:

$$\text{swingCapacity}(k) = \text{least of :}$$

$(\text{server.maxMinFlowRate}[k-1] - \sum_{i \in S_k} (\text{client.lookup}(i).\text{session.minimumFlowRate}))$
 and
 $(\text{server.maxFlowRate} - \sum_{i \in \text{ActiveClients}} (\text{client.lookup}(i).\text{session.minimumFlowRate}))$

The predicate expression invoked by step 1340 can now be written as follows:

$(\text{required_bandwidth} \leq \text{swingCapacity}(k) - \text{server.cac_flowSafetyMargin})$

This algorithm processes queued clients in band sequence, and within every band in FIFO. If the predicate evaluates to true the client is activated. Otherwise two possible cases are considered: 1) under the *FirstFit* embodiment the procedure continues scanning clients to the end of the banded queue, activating clients whose requirements can be met; 2) under the *FIFO* embodiment, the procedure ends with the first candidate client whose requirements cannot be met. Many other variations on these two embodiments might also be considered.

4.5.4 Anticipating Scheduled Variations in Network Capacity (Maximal Charge)

4.5.4.1 Overview

The procedure applicable to optimization of delivery charges is obtained by blending elements of the **CAC method depicted in Figure 13** into the method depicted in figure 11, which applies without change. To understand how this might work it may be useful to visualize a version of figure 10 stratified along its length in the manner of figure 8. As the maximum flow level undergoes a step change so too do the widths of its constituent strata in equal proportion.

4.5.4.1 Procedure

As previously mentioned, the method of **CAC in section 4.5.3** (figure 11) applies to this circumstance also, provided we alter the definition of two routines, (17) and (18), upon which that procedure relies, yielding (20) and (21), and adopt the banded queue organization outlined in the previous section.

The server swing capacity at a future time t is computed according to the capacity and worst-case client flows at time t .

(20) $\text{swingCapacity}(k, t) = \text{minimum of } ($
 $(\text{server.maxFlowRate}(t) * (\text{server.maxMinFlowRate}[k-1] / \text{server.maxMinFlowRate}[N-1]) -$
 $\sum_{i \in S_k \ \& \ (t_{\text{end}}(i) > t)} (\text{client.lookup}(i).\text{session.minFlowRate})) ,$
 $(\text{server.maxFlowRate} -$
 $\sum_{i \in \text{ActiveClients} \ \& \ (t_{\text{end}}(i) > t)} (\text{client.lookup}(i).\text{session.minFlowRate})))$

Finally, we define a predicate that tests whether a prospective customer will cause swing capacity to be exceeded at some time t , as follows:

(21) $\text{boolean client_fits}(i, t) \{$
 $k = \text{client.lookup}(i).\text{costOfService};$
 $\text{if}(\text{client.lookup}(i).\text{contentSelection.averagePlayRate} \leq$
 $\text{swingCapacity}(k, t) - \text{server.cac_flowSafetyMargin})$

```
        return true;

    else return false;

}
```

A method for call/connection acceptance and flow modulation for network delivery of video/audio programming is thus provided. Although several embodiments have been illustrated and described, it will be apparent to those skilled in the art that various changes and modifications may be made without departing from the spirit of the invention or the scope of the claims.

65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99

What is claimed is:

1. A method for optimal multimedia content delivery over networks from server to client comprising the steps of:

Delineating a state variable that represents the data rate to each client;

Delineating a set of requirements which represent the time-varying constraints on the data rate of said multimedia content; given by:

- (1) The total data rate for all clients does not exceed the maximum throughput of the server or network, whichever is least;
- (2) The data rate from server to client does not exceed the maximum data rate for the client;
- (3) The data rate to the client will never overflow the client buffer
- (4) The server will never underflow
- (6) The data rate from the server will never be less than the client's minimum data rate, which is a non-increasing function of time obtained by dividing the content not yet delivered by the remaining play time

Delineating a cost function which represents the value of a proposed solution

Performing periodic computations to solve said inequalities to obtain the state value that maximizes said cost function.

2. A method as in claim 1 further comprising:

said requirements further comprise:

The current maximum client data rate is given by the minimum of:

- The stored initial maximum client data rate;
- The data rate required to fill the remaining client buffer during the current of said periodic computations;
- The data rate required to complete the delivery of said multimedia content.
- The client data rate never exceeds said current maximum client data rate

Whereby said current maximum client data rate is periodically recomputed to maintain an optimal solution over a give period of time.

3. A method as in claim 2 further comprising:

Said cost function represents maximal throughput and is given by the sum of said client data rates for all active clients.

4. A method as in claim 2 further comprising:

Said cost function represents maximal charge and is given by the sum for all active clients of: said client data rates times the client's cost of service..

5. A method as in claim 3 for bandwidth allocation for delivery of multimedia data from server to one or more clients over a network, comprising the steps of:

Determining the maximum flow rate and minimum flow rate for each client;
Determining the flow rate range for each client as given by the difference between said maximum flow rate and said minimum flow rate;
Initializing current flow rate for each client as said minimum flow rate and summing said flow rate into total server flow rate; and
Allocating remaining server bandwidth to remaining clients until they each saturate or no bandwidth remains.

6. A method as in claim 5 wherein said step of allocating remaining server bandwidth to remaining clients is done fairly by a procedure that comprises the steps of:

Sorting the list of clients according to said flow rate range; and

Determining equally-allocated remaining server bandwidth if allocated evenly to all remaining unprocessed clients and;
Determining the range of remaining client bandwidth as given by the difference between said maximum flow rate and said minimum flow rate;
Determining saturation by comparing said equally-allocated remaining server bandwidth and said range of remaining client bandwidth, and allocating the lesser of these two amounts to each remaining client flow rate;

Whereby allocating flow to remaining clients based upon the sorted client range flow rates and determining allocation of remaining server bandwidth based upon a comparison of saturation of server versus saturation of each client maximizes allocation of total bandwidth for maximal flow rate to maximum number of clients.

7. A method as in claim 4 for bandwidth allocation for delivery of multimedia data from server to one or more clients over a network, comprising the steps of:

Determining the maximum flow rate and minimum flow rate for each client;
Determining the flow rate range for each client as given by the difference between said maximum flow rate and said minimum flow rate;
Sorting the list of clients according to said flow rate range;
Initializing current flow rate for each client as said minimum flow rate and summing said flow rate into total server flow rate;
Allocating remaining server bandwidth to remaining clients such that lower paying clients receive bandwidth only if higher paying ones are saturated.

8. A method as in claim 7 wherein said step of allocating remaining server bandwidth to remaining clients comprises the steps of:

For each remaining unprocessed client:

Determining equally-allocated remaining server bandwidth if allocated evenly to all remaining unprocessed clients and;
Determining the range of remaining client bandwidth as given by the difference between said maximum flow rate and said minimum flow rate;

Determining saturation by comparing said equally-allocated remaining server bandwidth and said range of remaining client bandwidth, and allocating the lesser of these two amounts to each remaining client flow rate;

Whereby allocating flow to remaining clients based upon the sorted client range flow rates and determining allocation of remaining server bandwidth based upon a comparison of saturation of server versus saturation of each client maximizes allocation of total bandwidth for maximal flow rate to maximum number of clients.

9. A method for connection acceptance control for delivery of multimedia data from server to one or more clients over a network, comprising the steps of:

Determining the server swing capacity given by the difference between the total server bandwidth and the sum of the minimum flow rate for each client;

Allocating server bandwidth for each prospective client which will fit without server bandwidth saturation, as determined by comparing the average data play rate of each prospective client with the remaining bandwidth available to the server.

10. A method as in claim 9 wherein said remaining bandwidth available to the server is given by said server swing capacity.

11. A method as in claim 10 wherein said remaining bandwidth available to the server is given by said server swing capacity less a server flow safety margin, thereby allowing server capacity to be subsequently lowered by up to the margin without requiring load shedding, and without affecting client sessions in process.

12. A method as in claim 9 wherein said step of allocating server bandwidth for each prospective client which will fit without server bandwidth saturation is further comprised of:

Allocating server bandwidth to each prospective client sequentially until a prospective client is located in which said average data play rate exceeds said sever swing capacity.

13. A method as in claim 9 wherein said step of allocating server bandwidth for each client which will fit without server bandwidth saturation is further comprised of:

Allocating server bandwidth to each prospective client sequentially for each client which can be activated without server bandwidth saturation.

14. A method for bandwidth allocation for delivery of multimedia data from server to one or more clients over a network, comprising the steps of:

Storing a sequence of data representing scheduled bandwidth changes for the server;

Determining the maximum flow rate and minimum flow rate for each client at the present time;

Determining the flow rate range for each client as given by the difference between said maximum flow rate and said minimum flow rate;

Sorting the list of clients according to said flow rate range;

Initializing current flow rate for each client as said minimum flow rate and summing said flow rate into total server flow rate; and
Allocating remaining server bandwidth to remaining clients.

Abstract

A method is disclosed for call and/or connection acceptance control and the optimal delivery of multimedia (audio/video) data over networks. This method involves the establishment and monitoring of certain criteria which may be used to maximize the number of simultaneous clients without sacrificing quality-of-service for already-connected clients. Methods are disclosed for maximizing total throughput as well as maximum charge models for different levels of service. The disclosed methods solve these optimization problems by expanding on linear-program techniques in manners geared towards multimedia content delivery over networks and many variations suitable for varying business models are disclosed.




Figure 1

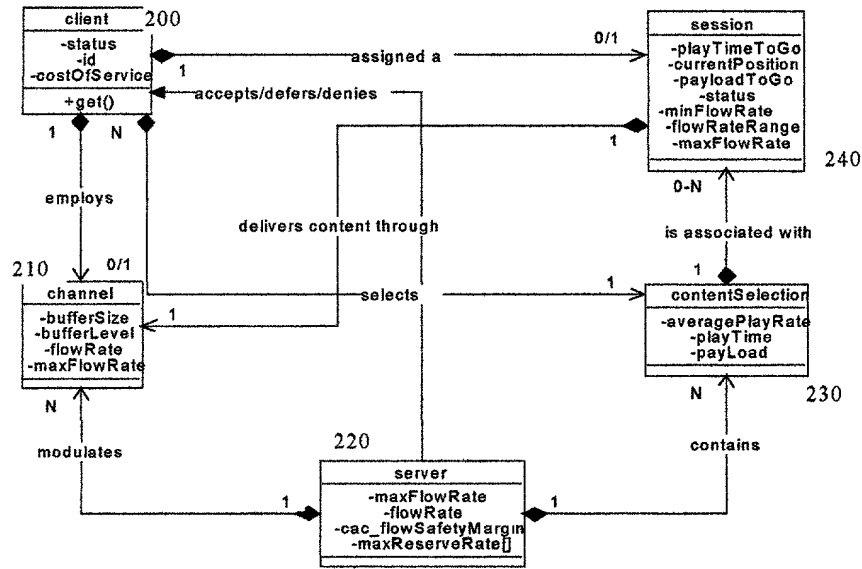


Figure 2




Figure 3




Figure 4

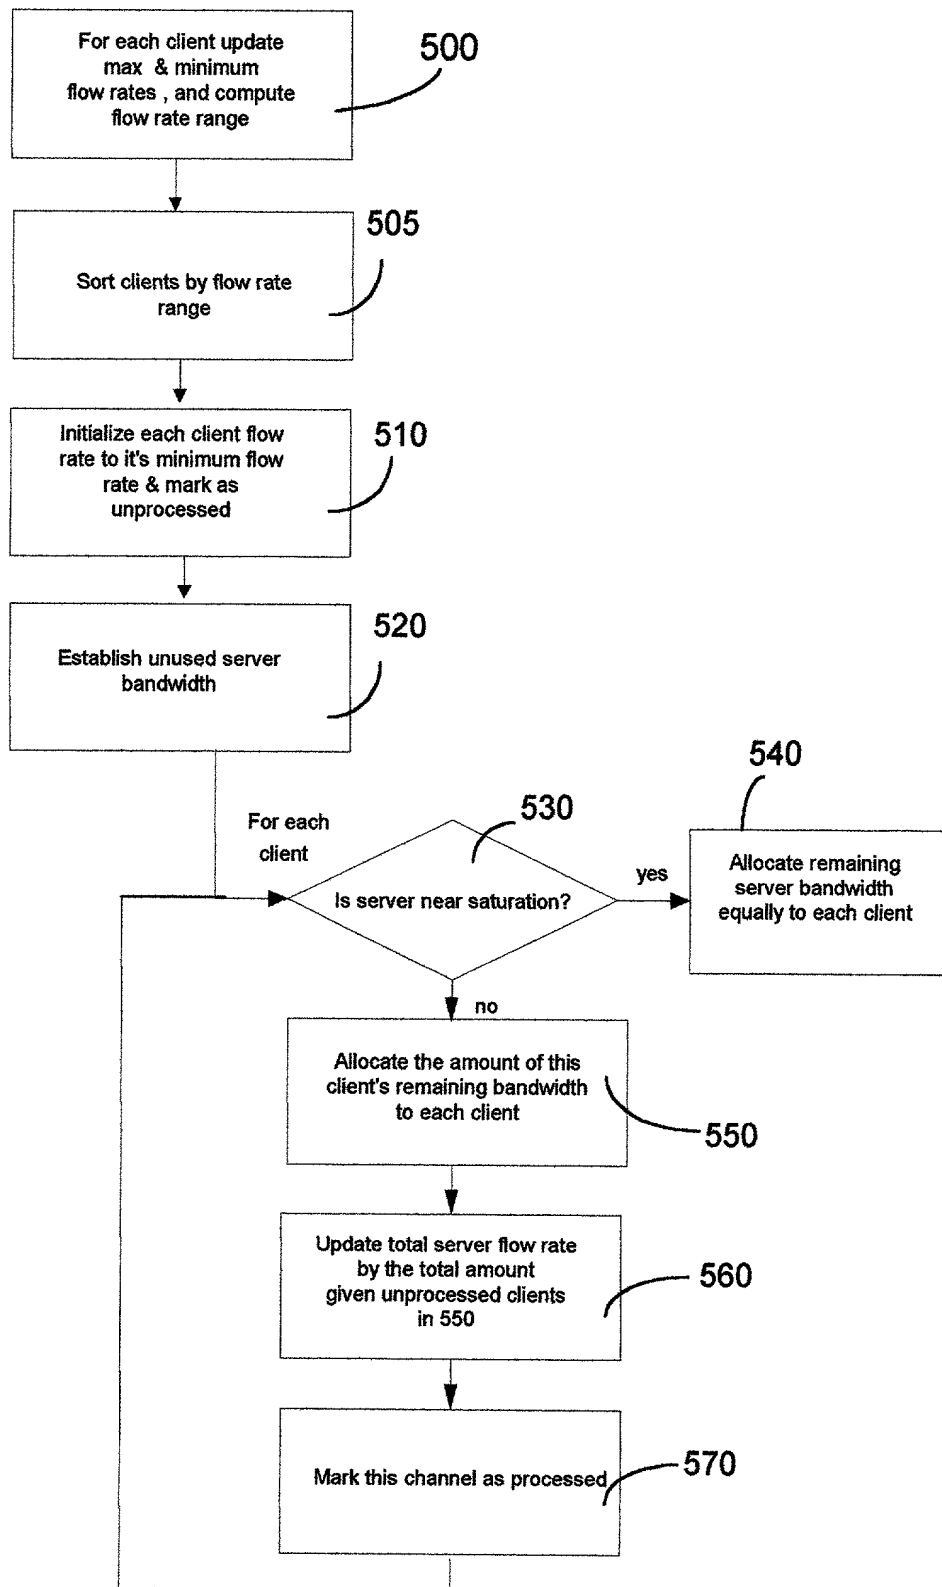


Figure 5

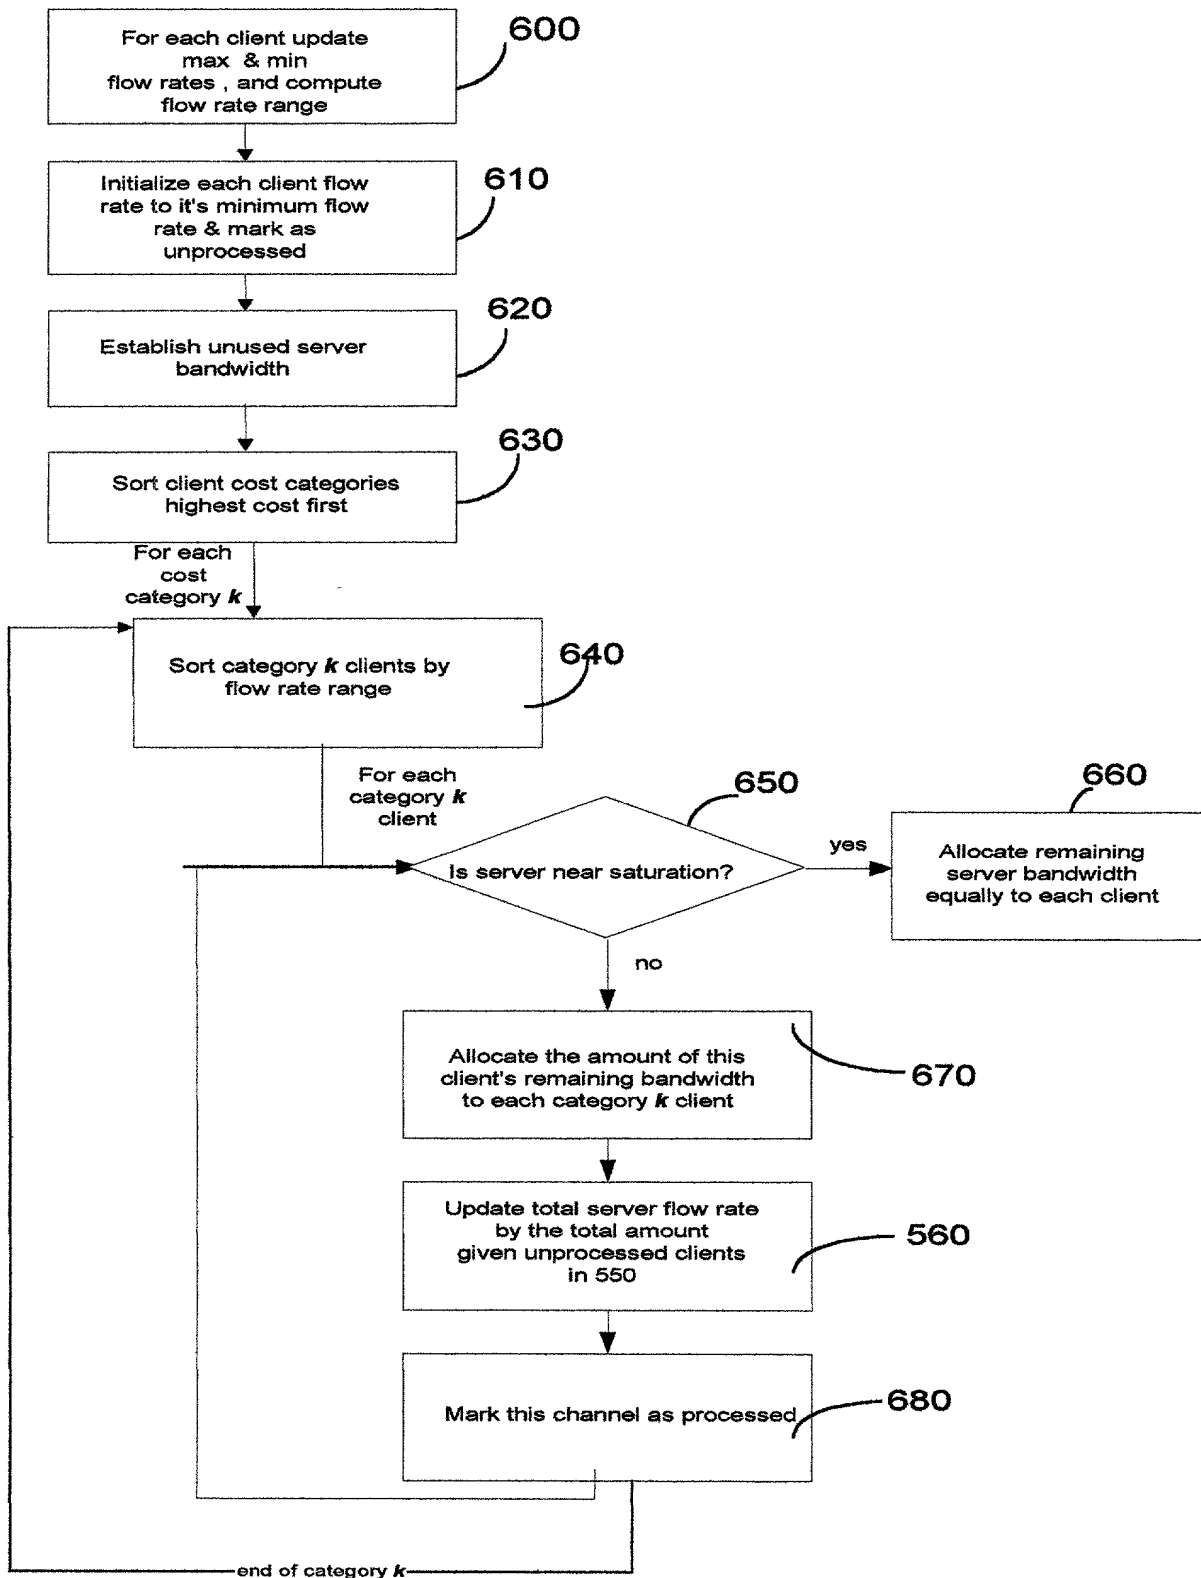


Figure 6




Figure 7




Figure 8

```

graph TD
    900[900 Compute server swing capacity] --> 920[920 compute client bandwidth requirement]
    920 --> 940{940 client fits}
    940 -- No (FIFO) --> 920
    940 -- yes --> 950[950 remove client from queue;  
Activate client & add to set of active clients]
    950 --> 960[960 decrement server swing capacity by client requirement]
    960 --> 920

```

name 33 of 37




Figure 10

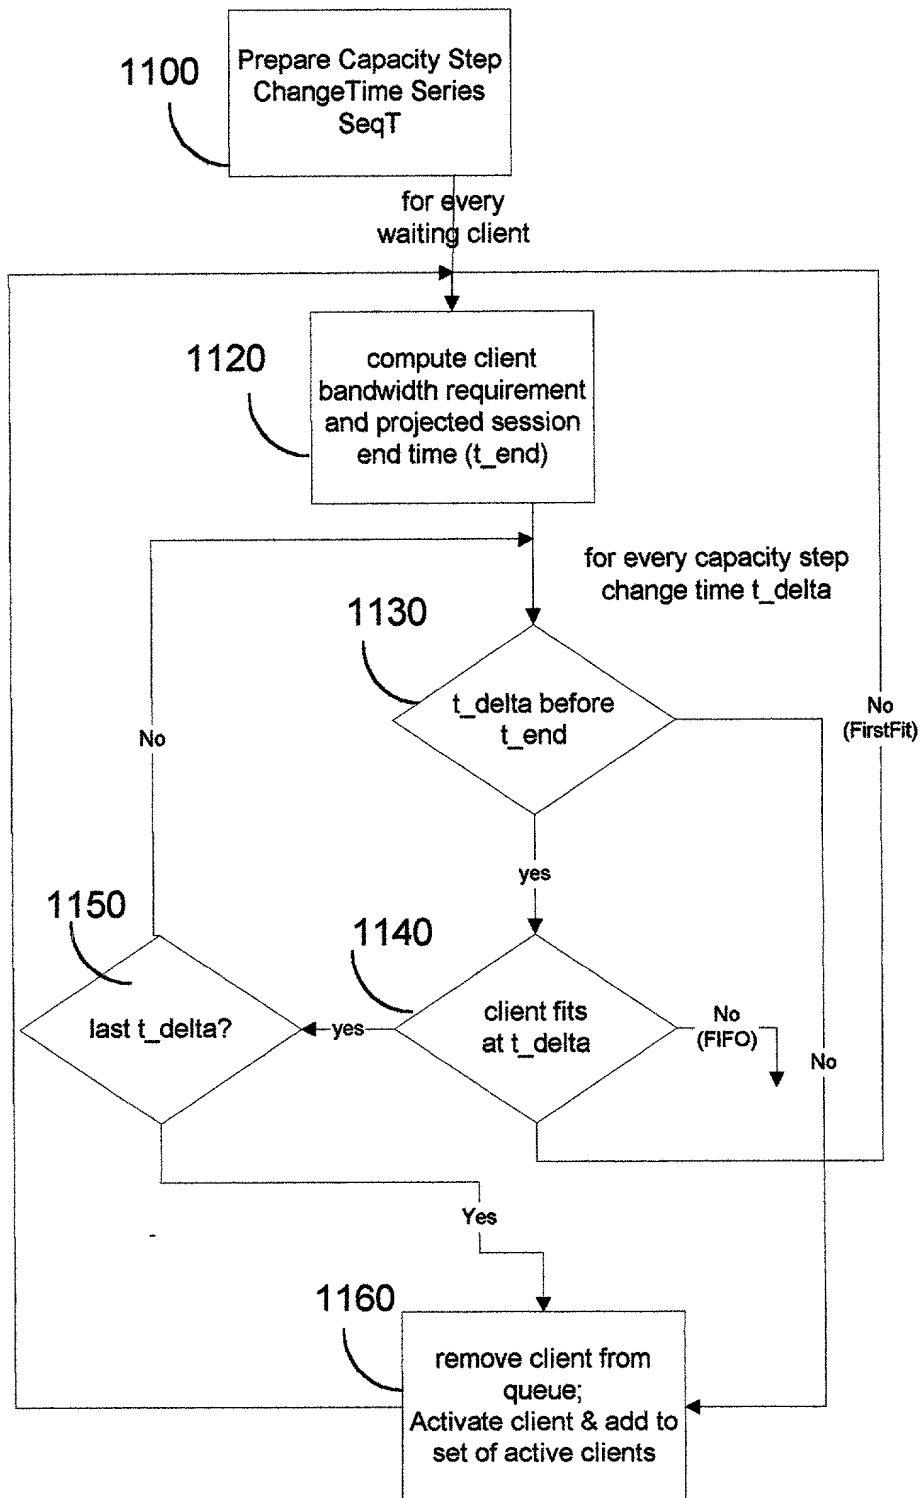


Figure 11




Figure 12




Figure 13

Please type a plus sign (+) inside this box → ☐

PTO/SB/01 (12-97)

Approved for use through 9/30/00. OMB 0651-0032
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION FOR UTILITY OR DESIGN PATENT APPLICATION (37 CFR 1.63) <input type="checkbox"/> Declaration Submitted with Initial Filing OR <input type="checkbox"/> Declaration Submitted after Initial Filing (surcharge (37 CFR 1.16 (e)) required)	Attorney Docket Number	BW02
	First Named Inventor	Arthur Allen
	COMPLETE IF KNOWN	
	Application Number	/
	Filing Date	6/25/99
	Group Art Unit	
	Examiner Name	

As a below named inventor, I hereby declare that:

My residence, post office address, and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

Method for Connection Acceptance Control and Optimal Multimedia Content Delivery Over Networks

the specification of which (Title of the Invention)

☒ is attached hereto
OR
☐ was filed on (MM/DD/YYYY) [] as United States Application Number or PCT International Application Number [] and was amended on (MM/DD/YYYY) [] (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR 1.56.

I hereby claim foreign priority benefits under 35 U.S.C. 119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate, or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or of any PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached?	
				YES	NO
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

☐ Additional foreign application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

I hereby claim the benefit under 35 U.S.C. 119(e) of any United States provisional application(s) listed below.

Application Number(s)	Filing Date (MM/DD/YYYY)
60/108,777	11/17/98

☐ Additional provisional application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

[Page 1 of 2]

Burden Hour Statement: This form is estimated to take 0.4 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Washington, DC 20231.

Please type a plus sign (+) inside this box → ☐

PTO/SB/01 (12-97)
Approved for use through 9/30/00. OMB 0651-0032
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION — Utility or Design Patent Application

I hereby claim the benefit under 35 U.S.C. 120 of any United States application(s), or 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application.

U.S. Parent Application or PCT Parent
Number

Parent Filing Date
(MM/DD/YYYY)

Parent Patent Number
(if applicable)

☐ Additional U.S. or PCT international application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

As a named inventor, I hereby appoint the following registered practitioner(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

☐ Customer Number

OR

☒ Registered practitioner(s) name/registration number listed below

Place Customer
Number Bar Code
Label here

Name	Registration Number	Name	Registration Number
Earl Mincer	36,396		

☐ Additional registered practitioner(s) named on supplemental Registered Practitioner Information sheet PTO/SB/02C attached hereto.

Direct all correspondence to: ☐ Customer Number
or Bar Code Label

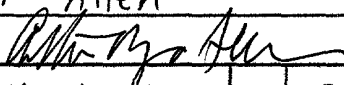
OR ☒ Correspondence address below

Name	Earl Mincer				
Address	4455 E Camelback				
Address	Suite E160				
City	Phoenix	State	AZ	ZIP	85018
Country	USA	Telephone	602 840-8411	Fax	

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Name of Sole or First Inventor:

☐ A petition has been filed for this unsigned inventor

Given Name (first and middle (if any))		Family Name or Surname					
Arthur Allen		Allen					
Inventor's Signature			Date				
Residence: City	Mountain View	State	CA	Country	USA	Citizenship	US
Post Office Address	1322 Isabelle Ave.						
Post Office Address							
City	Mountain View	State	CA	ZIP	94040	Country	USA

☐ Additional inventors are being named on the supplemental Additional Inventor(s) sheet(s) PTO/SB/02A attached hereto